

Exhibit 3

U.S. Patent No. 8,190,610

Charles Schwab Corporation





Claim Chart for Representative Claim 1

OVERVIEW



U.S. Patent No. 8,190,610

 US008190610B2	
(12) United States Patent Dasdan et al.	(10) Patent No.: US 8,190,610 B2 (45) Date of Patent: May 29, 2012
(54) MAPREDUCE FOR DISTRIBUTED DATABASE PROCESSING	
(75) Inventors: Ali Dasdan, San Jose, CA (US); Hung-Chih Yang, Sunnyvale, CA (US); Ruey-Lung Hsiao, Los Angeles, CA (US)	
(73) Assignee: Yahoo! Inc., Sunnyvale, CA (US)	
(*) Notice: Subject to any disclaimer, the term of this patent is extended or adjusted under 35 U.S.C. 154(b) by 1105 days.	
(21) Appl. No.: 11/539,090	
(22) Filed: Oct. 5, 2006	
(65) Prior Publication Data US 2008/0086442 A1 Apr. 10, 2008	
(51) Int. Cl. G06F 17/30 (2006.01)	
(52) U.S. Cl. 707/737, 707/968	
(58) Field of Classification Search 707/1-3, 707/737, 968 See application file for complete search history.	
(56) References Cited U.S. PATENT DOCUMENTS 6,158,644 A * 12/2000 Tibbatts 717/100 6,341,230 B1 * 1/2002 Burroughs et al. 707/737	
(57) OTHER PUBLICATIONS Jeffrey Dean and Sanjay Ghemawat, "MapReduce: Simplified Data Processing on Large Clusters", USENIX Association OSDI '04: 6th Symposium on Operating Systems Design and Implementation, Dec. 6-8, 2004, pp. 137-148. * cited by examiner Primary Examiner — Khanh Pham (74) Attorney, Agent or Firm — Weaver Austin Villeneuve & Sampson LLP	
(57) ABSTRACT An input data set is treated as a plurality of grouped sets of key/value pairs, which enhances the utility of the MapReduce programming methodology. By utilizing such a grouping, map processing can be carried out independently on two or more related but possibly heterogeneous datasets (e.g., related by being characterized by a common primary key). The intermediate results of the map processing (key/value pairs) for a particular key can be processed together in a single reduce function by applying a different iterator to intermediate values for each group. Different iterators can be arranged inside reduce functions in ways however desired.	
46 Claims, 5 Drawing Sheets	



Title: MAPREDUCE FOR DISTRIBUTED DATABASE PROCESSING

Priority Date: October 5, 2006

Filing Date: October 5, 2006

Issue Date: May 29, 2012

Expiration Date: October 14, 2029

(57) ABSTRACT

An input data set is treated as a plurality of grouped sets of key/value pairs, which enhances the utility of the MapReduce programming methodology. By utilizing such a grouping, map processing can be carried out independently on two or more related but possibly heterogeneous datasets (e.g., related by being characterized by a common primary key). The intermediate results of the map processing (key/value pairs) for a particular key can be processed together in a single reduce function by applying a different iterator to intermediate values for each group. Different iterators can be arranged inside reduce functions in ways however desired.

Representative Claim 1

A method of processing data of a data set over a distributed system, wherein the data set comprises a plurality of data groups, the method comprising:

partitioning the data of each one of the data groups into a plurality of data partitions that each have a plurality of key-value pairs and providing each data partition to a selected one of a plurality of mapping functions that are each user-configurable to independently output a plurality of lists of values for each of a set of keys found in such map function's corresponding data partition to form corresponding intermediate data for that data group and identifiable to that data group, wherein the data of a first data group has a different schema than the data of a second data group and the data of the first data group is mapped differently than the data of the second data group so that different lists of values are output for the corresponding different intermediate data, wherein the different schema and corresponding different intermediate data have a key in common; and

reducing the intermediate data for the data groups to at least one output data group, including processing the intermediate data for each data group in a manner that is defined to correspond to that data group, so as to result in a merging of the corresponding different intermediate data based on the key in common,

wherein the mapping and reducing operations are performed by a distributed system.

CLAIM CHART



References Cited

Upon information and belief, Charles Schwab has and continues to use a system called Big Data Analytics (hereinafter, the “Charles Schwab system”), which is an analytics platform built on technology from Datameer layered on top of an Apache Hadoop framework that includes the elements in the systems described in and/or practices the steps of the methods described in the claims of U.S. Patent No. 8,190,610 (the “610 Patent”).

The following chart presents R2 Solutions’ analysis of the Charles Schwab system based on publicly available information.

The citations in the chart refer to the following publicly available documents, which are incorporated by reference as if fully set forth herein:

- [1] Sachin P. Bappalige, “An introduction to Apache Hadoop for big data,” opensource.com (Aug. 26, 2014), *available at* <https://opensource.com/life/14/8/intro-apache-hadoop-big-data> (“**An introduction to Apache Hadoop for big data**”)
- [2] MapReduce Tutorial, *available at* <http://hadoop.apache.org/docs/current/hadoop-mapreduce-client/hadoop-mapreduce-client-core/MapReduceTutorial.html> (“**MapReduce Tutorial**”)
- [3] Interface Reducer<K2,V2,K3,V3>, *available at* <https://hadoop.apache.org/docs/r2.8.0/api/org/apache/hadoop/mapred/Reducer.html> (“**Interface Reducer**”)
- [4] ASF Infrabot, “HadoopMapReduce,” Dashboard (July 9, 2019), *available at* <https://cwiki.apache.org/confluence/display/HADOOP2/HadoopMapReduce> (“**HadoopMapReduce**”)
- [5] Tom Davenport, “Feeding a Data-Hungry Organization at Charles Schwab,” Forbes (Jun. 5, 2018), *available at* <https://www.forbes.com/sites/tomdavenport/2018/06/05/feeding-a-data-hungry-organization-at-charles-schwab/?sh=541cbb3977bb> (“**Feeding a Data-Hungry Organization**”)
- [6] “Charles Schwab Careers: Senior Data Engineer,” *available at* <https://career-schwab.icims.com/jobs/64871/senior-data-engineer---etl-lead/job?mobile=false&width=1230&height=500&bga=true&needsRedirect=false&jan1offset=-360&jun1offset=-300>

References Cited (continued)

- [7] *“Charles Schwab Careers: Senior Software Engineer,”* available at <https://jobs.schwabjobs.com/job/san-francisco/senior-software-engineer/33727/17647572>
- [8] *“Internship at Charles Schwab – Summer 2019,”* UT Dallas, available at <https://jindal.utdallas.edu/som/internship-stories/1087>
- [9] *LinkedIn Profile: Guarav Singh,* available at <https://www.linkedin.com/in/gaurav-singh-82812614> (“**LinkedIn Profile**”)
- [10] Hive 2.1.1 API – Class Schema, available at <https://hive.apache.org/javadocs/r2.1.1/api/org/apache/hadoop/hive/metastore/api/Schema.html> (“**Class Schema**”)
- [11] Confluence Administrator, “Hive Tutorial,” Dashboard (updated by Owen O’Malley, Mar. 5, 2019), available at <https://cwiki.apache.org/confluence/display/Hive/Tutorial> (“**Hive Tutorial**”)
- [12] Confluence Administrator, “LanguageManual Joins,” Dashboard (updated by Dudu Markovitz, June 5, 2017), available at <https://cwiki.apache.org/confluence/display/Hive/LanguageManual+Joins> (“**LanguageManual Joins**”)
- [13] Tanel Poder, “Speed Up Your Queries with Hive LLAP Engine on Hadoop or in the Cloud,” Gluent, available at <https://www.slideshare.net/gluent/speed-up-your-queries-with-hive-llap-engine-on-hadoop-or-in-the-cloud> (“**Speed Up Your Queries with Hive LLAP Engine**”)
- [14] “Apache Hive,” Wikipedia, available at https://en.wikipedia.org/wiki/Apache_Hive (“**Apache Hive**”)
- [15] Petar Zecevic & Marko Bonaci, “Spark in Action” (Manning Publications Nov. 2016), available at https://learning.oreilly.com/library/view/spark-in-action/9781617292606/kindle_split_014.html (“**Spark in Action**”)
- [16] “Apache Spark FAQ,” available at <https://spark.apache.org/faq.html> (“**Apache Spark FAQ**”)

References Cited (continued)

- [17] Apache Spark 3.0.1, “RDD Programming Guide,” *available at* <http://spark.apache.org/docs/latest/rdd-programming-guide.html> (“**RDD Programming Guide**”)
- [18] “Learning Spark,” O’Reilly Media, Inc. (February 2015), *available at* <https://learning.oreilly.com/library/view/learning-spark/9781449359034/ch04.html#chap-pair-RDDS> (“**Learning Spark**”)
- [19] “Spark RDD Transformations with examples,” *available at* <https://sparkbyexamples.com/apache-spark-rdd/spark-rdd-transformations/> (“**Spark RDD Transformations**”)
- [20] Shrey Mehrotra & Akash Grade, “Apache Quick Start Guide” (Packt Publishing, Jan. 2019), *available at* <https://learning.oreilly.com/library/view/apache-spark-quick/9781789349108/0ee1a5e2-09d0-49f0-99f5-9dee8336258d.xhtml> (“**Apache Spark Quick Start Guide**”)
- [21] “Spark SQL Shuffle Partitions,” *available at* <https://sparkbyexamples.com/spark/spark-shuffle-partitions/> (“**Spark SQL Shuffle Partitions**”)
- [22] “Apache Spark Map vs. FlatMap Operation,” DataFlair, *available at* <https://data-flair.training/blogs/apache-spark-map-vs-flatmap/> (“**Apache Spark Map vs. FlatMap**”)
- [23] Confluence Administrator, “LanguageManual Transform,” Dashboard (updated by Sushanth Sowmyan on Dec. 9, 2015), *available at* <https://cwiki.apache.org/confluence/display/Hive/LanguageManual+Transform> (“**LanguageManual Transform**”)

Claim Chart – Claim 1

Claim Feature	Evidence from Charles Schwab
<p>A method of processing data of a data set over a distributed system, wherein the data set comprises a plurality of data groups, the method comprising:</p>	<p>Charles Schwab infringes Claim 1 of the '610 Patent. The Charles Schwab system performs a method of processing data of a data set over a distributed system, wherein the data set comprises a plurality of data groups, the method comprising the steps discussed below:</p> <p>Charles Schwab uses an Apache Hadoop distributed file system that relies on MapReduce.</p> <p>[5] Feeding a Data-Hungry Organization:</p> <div data-bbox="747 784 1839 1240" style="border: 1px solid black; padding: 10px;"> <p>At about the same time, <u>Schwab felt that it needed to modernize its data environment</u> and Andrew Salesky was named head of Global Data. The Global Data Technology organization was also formed, reporting to the CIO. <u>This led to incorporating a Hadoop-based big data lake and other new technologies.</u> Salesky notes that Schwab, like any other organization wanting to do more with advanced analytics and AI, needed to make some “long tailed” investments in its data. In terms of data integration and master data management, the primary focus at Schwab has been on client data and developing consistent enterprise definitions of key data assets.</p> </div>


Claim Chart – Claim 1

Claim Feature	Evidence from Charles Schwab
A method of processing data of a data set over a distributed system, wherein the data set comprises a plurality of data groups, the method comprising:	<p>Upon information and belief, the Charles Schwab System uses Apache Hadoop, Hive, and Spark:</p> <p>[6] Charles Schwab Careers: Senior Data Engineer:</p> <p>Your Opportunity</p> <p>Charles Schwab & Co., Inc is currently seeking a seasoned ETL Lead with a passion for hands-on design/development and collaboration with our business partners. The ideal candidate is adept at using large data sets to find opportunities for product and process optimization to test the effectiveness of different courses of action. The ETL Lead must have deep experience in Enterprise Data Warehouse, Data Mart design, development and end to end execution of data solutions. Prior experience in developing design patterns for Big Data and Cloud solutions is a big plus.</p> <p>This role will require hands on development with wide range of technical skills which include data analysis, ETL (Talend / Informatica, Unix, Big Data Technologies) and modeling skills (SQL, ERwin).</p> <p>What you are good at</p> <p>This position is part of the Global Data Technology (GDT) organization that governs the strategy and implementation of the enterprise data warehouse and emerging data platforms.</p> <p><u>The ETL Lead will be designing, building and supporting data processing pipelines to transform data using Hadoop technologies.</u> You'll be designing schemas, data models and data architecture for <u>Hadoop</u> and HBase environments. You'll be implementing data flow scripts using Unix / <u>Hive QL</u> / Pig scripting. <u>You'll be designing, building data assets in MapR-DB (HBASE), and HIVE.</u> You'll be developing and executing quality assurance and test scripts. You'll be work with product owners and business analysts to understand business requirements and use cases to design solutions. You'll have the opportunity to grow in responsibility, work on exciting and challenging projects, train on emerging technologies and help set the future of the Data Solution Delivery team. You'll lead investigation and resolution efforts for critical/high impact problems, defects and incidents. Provides technical guidance to team members.</p>

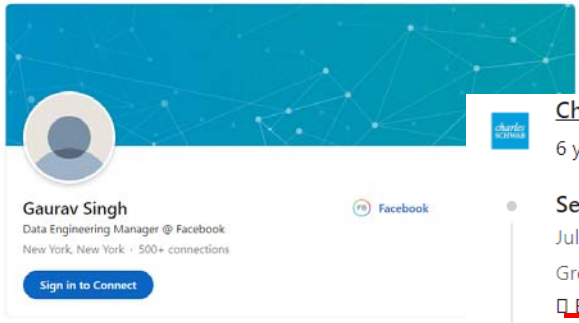
Claim Chart – Claim 1

Claim Feature	Evidence from Charles Schwab
<p>A method of processing data of a data set over a distributed system, wherein the data set comprises a plurality of data groups, the method comprising:</p>	<p>Upon information and belief, the Charles Schwab System uses Apache Hadoop, Hive, and Spark:</p> <p>[7] Charles Schwab Careers: Senior Software Engineer:</p> <div data-bbox="653 743 1980 1114"> <p>Responsibilities</p> <ul style="list-style-type: none"> • <u>Build reliable data pipelines</u> using Python and Scala (<u>with Apache Spark</u>) to ingest, process and transform data • Enable easy interaction with processed data by exposing REST-like endpoints • Write algorithms that interact with the processed data to transform datasets into portfolios based on various parameters • Add functionality to a web application to wrap the backend services and make them easier to use • Write extensive and reliable tests at all parts of the test pyramid • Continuously make improvements and optimizations so that we can build portfolios using large amounts of data in seconds • Write automated tests for all code to ensure the highest quality product • Contribute to architectural decisions related to the platform • Build infrastructure required to deploy the code in the cloud using Terraform </div>

Claim Chart – Claim 1

Claim Feature	Evidence from Charles Schwab
<p>A method of processing data of a data set over a distributed system, wherein the data set comprises a plurality of data groups, the method comprising:</p>	<p>Upon information and belief, the Charles Schwab System uses Apache Hadoop, Hive, and Spark:</p> <p>[8] Internship at Charles Schwab – Summer 2019:</p>  <p>I am involved in the 2 major projects and a small project. The major projects I dealt with were connected to the Digital Accelerators namely DARE and ONCE. In this project I had to make and automate the Web Analytics DARE and ONCE dashboards in Tableau. Initially the web dashboards used to get updated with the manual intervention by pulling out the Adobe Analytics Data and storing it into excel and then connecting Excel to Tableau. For automating this process, <u>the Adobe Analytics Data was loaded into the Big Data Environment where the system for the project was Hadoop and I used the Nexus Query Chameleon to query the big data with the help of Apache Hive – HQL.</u> Results of the queries were stored in the IDW database. There after the database was connected to tableau and I prepared the Dashboards. Automation of this whole process is done by running some of the UNIX jobs.</p>

Claim Chart – Claim 1

Claim Feature	Evidence from Charles Schwab
<p>A method of processing data of a data set over a distributed system, wherein the data set comprises a plurality of data groups, the method comprising:</p>	<p>Upon information and belief, the Charles Schwab System uses Apache Hadoop, Hive, and Spark:</p> <p>[9] LinkedIn Profile:</p>  <p>Charles Schwab 6 years 2 months</p> <p>Senior Manager, Data Engineering Jul 2016 - Jul 2019 · 3 years 1 month Greater Denver Area</p> <ul style="list-style-type: none"> □ <u>Building strong data pipeline using HDFS, Python, Hive, Spark, Kafka & Teradata.</u> □ Responsible for the development of a big data customer behavior hub to enable cross channel attribution and marketing optimization using HDFS, Hive, SPARK & Python. □ Responsible for Strategy, Design & Delivery functions aimed at creating a robust analytical data layer on Enterprise Data warehouse platform to help facilitate “one version of the truth” reporting across multiple product and delivery □ Manage a team 35+ engineers in diverse geographical locations and varied roles like Data Engineer, Business System Analyst, Scrum masters etc. □ Responsible for multi million project portfolio annually. <p>Show less ^</p>

Claim Chart – Claim 1

Claim Feature	Evidence from Charles Schwab
<p>A method of processing data of a data set over a distributed system, wherein the data set comprises a plurality of data groups, the method comprising:</p>	<p>Charles Schwab's Apache Hadoop MapReduce implementation is over a distributed system:</p> <p>[1] An introduction to Apache Hadoop for big data at 2:</p> <div style="border: 1px solid black; padding: 10px; margin: 10px 0;"> <p>The Apache Hadoop framework is composed of the following modules</p> <ol style="list-style-type: none"> 1. Hadoop Common: contains libraries and utilities needed by other Hadoop modules 2. Hadoop Distributed File System (HDFS): <u>a distributed file-system that stores data on the commodity machines, providing very high aggregate bandwidth across the cluster</u> 3. <u>Hadoop YARN</u>: a resource-management platform responsible for managing compute resources in clusters and using them for scheduling of users' applications 4. <u>Hadoop MapReduce</u>: a programming model for large scale data processing <p><u>All the modules in Hadoop are designed with a fundamental assumption that hardware failures (of individual machines, or racks of machines) are common and thus should be automatically handled in software by the framework.</u> Apache Hadoop's MapReduce</p> </div>

Claim Chart – Claim 1

Claim Feature	Evidence from Charles Schwab
A method of processing data of a data set over a distributed system , wherein the data set comprises a plurality of data groups, the method comprising:	<p>[14] Apache Hive:</p> <p><u>Apache Hive is a data warehouse software project built on top of Apache Hadoop for providing data query and analysis.</u>^[3] Hive gives an <u>SQL-like interface</u> to query data stored in various databases and file systems that integrate with Hadoop. Traditional SQL queries must be implemented in the <u>MapReduce Java API</u> to execute SQL applications and queries over distributed data. Hive provides the necessary SQL abstraction to integrate SQL-like queries (<u>HiveQL</u>) into the underlying Java without the need to implement queries in the low-level Java API. Since most data warehousing applications work with SQL-based querying languages, Hive aids portability of SQL-based applications to Hadoop.^[4] While initially developed by <u>Facebook</u>, Apache Hive is used and developed by other companies such as <u>Netflix</u> and the <u>Financial Industry Regulatory Authority (FINRA)</u>.^{[5][6]} Amazon maintains a software fork of Apache Hive included in <u>Amazon Elastic MapReduce</u> on <u>Amazon Web Services</u>.^[7]</p>

Claim Chart – Claim 1

Claim Feature	Evidence from Charles Schwab
<p>A method of processing data of a data set over a distributed system, wherein the data set comprises a plurality of data groups, the method comprising:</p>	<p>[16] Apache Spark FAQ at 1:</p> <div data-bbox="632 529 2001 695"> <p>How does Spark relate to Apache Hadoop?</p> <p>Spark is a fast and general processing engine compatible with Hadoop data. It can run in Hadoop clusters through YARN or Spark's standalone mode, and it can process data in HDFS, HBase, Cassandra, Hive, and any Hadoop InputFormat. It is designed to perform both batch processing (similar to MapReduce) and new workloads like streaming, interactive queries, and machine learning.</p> </div> <p>[17] RDD Programming Guide at 4, 2:</p> <div data-bbox="699 792 1944 1219"> <h3>Resilient Distributed Datasets (RDDs)</h3> <p>Spark revolves around the concept of a <i>resilient distributed dataset</i> (RDD), which is a fault-tolerant collection of elements that can be operated on in parallel. There are two ways to create RDDs: <u>parallelizing an existing collection in your driver program, or referencing a dataset in an external storage system, such as a shared filesystem, HDFS, HBase, or any data source offering a Hadoop InputFormat.</u></p> <h3>Overview</h3> <p>At a high level, every Spark application consists of a <i>driver program</i> that runs the user's main function and executes various <i>parallel operations</i> on a cluster. <u>The main abstraction Spark provides is a <i>resilient distributed dataset</i> (RDD), which is a collection of elements partitioned across the nodes of the cluster that can be operated on in parallel.</u> RDDs are created by starting with a file in the Hadoop file system (or any other Hadoop-supported file system), or an existing Scala collection in the driver program, and transforming it. Users may also ask Spark to <i>persist</i> an RDD in memory, allowing it to be reused efficiently across parallel operations. Finally, RDDs automatically recover from node failures.</p> </div>

Claim Chart – Claim 1

Claim Feature	Evidence from Charles Schwab
<p>partitioning the data of each one of the data groups into a plurality of data partitions that each have a plurality of key-value pairs and providing each data partition to a selected one of a plurality of mapping functions that are each user-configurable to independently output a plurality of lists of values for each of a set of keys found in such map function's corresponding data partition to form corresponding intermediate data for that data group and identifiable to that data group,</p>	<p>The method performed by the Charles Schwab system includes partitioning the data of each one of the data groups into a plurality of data partitions that each have a plurality of key-value pairs and providing each data partition to a selected one of a plurality of mapping functions that are each user-configurable to independently output a plurality of lists of values for each of a set of keys found in such map function's corresponding data partition to form corresponding intermediate data for that data group and identifiable to that data group.</p> <p>In Charles Schwab's Hadoop MapReduce framework, the input data is partitioned into independent chunks (data partitions), which are processed by Mapper functions:</p> <p>[2] MapReduce Tutorial at 2:</p> <div data-bbox="814 917 1667 1198" style="border: 1px solid black; padding: 10px;"> <p><u>A MapReduce job usually splits the input data-set into independent chunks which are processed by the map tasks in a completely parallel manner.</u> The framework sorts the outputs of the maps, which are then input to the reduce tasks. Typically both the input and the output of the job are stored in a file-system. The framework takes care of scheduling tasks, monitoring them and re-executes the failed tasks.</p> </div>

Claim Chart – Claim 1

Claim Feature	Evidence from Charles Schwab
<p>partitioning the data of each one of the data groups into a plurality of data partitions that each have a plurality of key-value pairs and providing each data partition to a selected one of a plurality of mapping functions that are each user-configurable to independently output a plurality of lists of values for each of a set of keys found in such map function's corresponding data partition to form corresponding intermediate data for that data group and identifiable to that data group,</p>	<p>The data partitions with key-value pairs are passed to user-configurable Mapper functions for further processing:</p> <p>[4] HadoopMapReduce at 1:</p> <div data-bbox="690 667 1915 1040" style="border: 1px solid black; padding: 10px;"> <p><u>As key-value pairs are read from the RecordReader they are passed to the configured Mapper. The user supplied Mapper does whatever it wants with the input pair and calls OutputCollector.collect with key-value pairs of its own choosing. The output it generates must use one key class and one value class. This is because the Map output will be written into a SequenceFile which has per-file type information and all the records must have the same type (use subclassing if you want to output different data structures). The Map input and output key-value pairs are not necessarily related typewise or in cardinality.</u></p> </div>

Claim Chart – Claim 1

Claim Feature	Evidence from Charles Schwab
<p>partitioning the data of each one of the data groups into a plurality of data partitions that each have a plurality of key-value pairs and providing each data partition to a selected one of a plurality of mapping functions that are each user-configurable to independently output a plurality of lists of values for each of a set of keys found in such map function's corresponding data partition to form corresponding intermediate data for that data group and identifiable to that data group,</p>	<p>The Mapper function maps the key-value pairs to form intermediate key-value pairs:</p> <p>[2] MapReduce Tutorial at 7:</p> <div data-bbox="648 699 1990 922" style="border: 1px solid black; padding: 10px; margin: 10px 0;"> <p>Mapper</p> <p>Mapper maps <u>input key/value pairs</u> to a set of <u>intermediate key/value pairs</u>.</p> <p><u>Maps are the individual tasks that transform input records into intermediate records.</u> The transformed intermediate records do not need to be of the same type as the input records. A given input pair may map to zero or many output pairs.</p> </div>

Claim Chart – Claim 1

Claim Feature	Evidence from Charles Schwab
<p>partitioning the data of each one of the data groups into a plurality of data partitions that each have a plurality of key-value pairs and providing each data partition to a selected one of a plurality of mapping functions that are each user-configurable to independently output a plurality of lists of values for each of a set of keys found in such map function's corresponding data partition to form corresponding intermediate data for that data group and identifiable to that data group,</p>	<p>With Spark, a data set is first partitioned into a plurality of elements and distributed across nodes. This distribution of elements is called a resilient distributed dataset (RDD):</p> <p>[17] RDD Programming Guide at 2, 4, 5:</p> <div data-bbox="699 656 1827 1203"> <p>Overview</p> <p>At a high level, every Spark application consists of a <i>driver program</i> that runs the user's main function and executes various <i>parallel operations</i> on a cluster. <u>The main abstraction Spark provides is a <i>resilient distributed dataset</i> (RDD), which is a collection of elements partitioned across the nodes of the cluster that can be operated on in parallel. RDDs are created by starting with a file in the Hadoop file system (or any other Hadoop-supported file system), or an existing Scala collection in the driver program, and transforming it. Users may also ask Spark to <i>persist</i> an RDD in memory, allowing it to be reused efficiently across parallel operations. Finally, RDDs automatically recover from node failures.</u></p> <p>Resilient Distributed Datasets (RDDs)</p> <p>Spark revolves around the concept of a <i>resilient distributed dataset</i> (RDD), which is a fault-tolerant collection of elements that can be operated on in parallel. There are two ways to create RDDs: <u>parallelizing an existing collection in your driver program, or referencing a dataset in an external storage system, such as a shared filesystem, HDFS, HBase, or any data source offering a Hadoop InputFormat.</u></p> <p>One important parameter for parallel collections is the number of <i>partitions</i> to cut the dataset into. Spark will run one task for each partition of the cluster. Typically you want 2-4 partitions for each CPU in your cluster. Normally, Spark tries to set the number of partitions automatically based on your cluster. However, you can also set it manually by passing it as a second parameter to <code>parallelize</code> (e.g. <code>sc.parallelize(data, 10)</code>). Note: some places in the code use the term <i>slices</i> (a synonym for partitions) to maintain backward compatibility.</p> </div>

Claim Chart – Claim 1

Claim Feature

partitioning the data of each one of the data groups into a plurality of data partitions that each have a plurality of key-value pairs and providing each data partition to a selected one of a plurality of mapping functions that are each user-configurable to independently output a plurality of lists of values for each of a set of keys found in such map function's corresponding data partition to form corresponding intermediate data for that data group and identifiable to that data group,

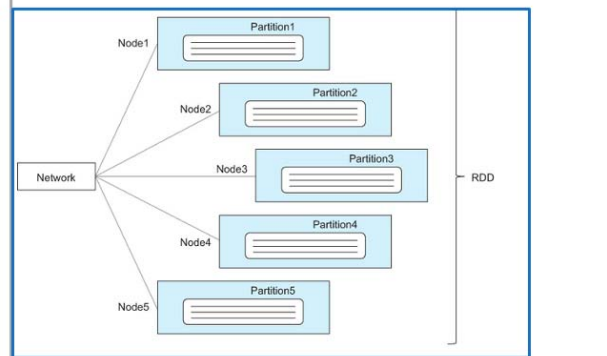
Evidence from Charles Schwab

[15]

Spark in Action:

Each part (piece or slice) of an RDD is called a *partition*.^[2] When you load a text file from your local filesystem into Spark, for example, the file's contents are split into partitions, which are evenly distributed to nodes in a cluster. More than one partition may end up on the same node. The sum of all those partitions forms your RDD. This is where the word *distributed* in *resilient distributed dataset* comes from. Figure 4.1 shows the distribution of lines of a text file loaded into an RDD in a five-node cluster. The original file had 15 lines of text, so each RDD partition was formed with 3 lines of text. Each RDD maintains a list of its partitions and an optional list of preferred locations for computing the partitions.

Figure 4.1. Simplified look at partitions of an RDD in a five-node cluster. The RDD was created by loading a text file using the `textFile` method of `SparkContext`. The loaded text file had 15 lines of text, so each partition was formed with 3 lines of text.



Claim Chart – Claim 1

Claim Feature	Evidence from Deezer
<p>partitioning the data of each one of the data groups into a plurality of data partitions that each have a plurality of key-value pairs and providing each data partition to a selected one of a plurality of mapping functions that are each user-configurable to independently output a plurality of lists of values for each of a set of keys found in such map function's corresponding data partition to form corresponding intermediate data for that data group and identifiable to that data group,</p>	<p>One form of the RDDs is a pair RDD, which includes key/value pairs:</p> <p>[18] Learning Spark:</p> <p>Spark provides special operations on <u>RDDs containing key/value pairs</u>. These RDDs are called <u>pair RDDs</u>. Pair RDDs are a useful building block in many programs, as they expose operations that allow you to act on each key in parallel or regroup data across the network. For example, pair RDDs have a <code>reduceByKey()</code> method that can aggregate data separately for each key, and a <code>join()</code> method that can merge two RDDs together by grouping elements with the same key. It is common to extract fields from an RDD (representing, for instance, an event time, customer ID, or other identifier) and use those fields as keys in pair RDD operations.</p> <p>[15] Spark in Action:</p> <p>4.1. WORKING WITH PAIR RDDS</p> <p>Storing data as key-value pairs offers a simple, general, and extensible data model because each <u>key-value pair can be stored independently</u> and it's easy to add new types of keys and new types of values. This extensibility and simplicity have made this practice fundamental to several frameworks and applications. For example, many popular caching systems and NoSQL databases, such as memcached and Redis, are key-value stores. Hadoop's MapReduce also operates on key-value pairs (as you can see in appendix A).</p> <p>Keys and values can be simple types such as integers and strings, as well as more-complex data structures. Data structures traditionally used to represent key-value pairs are <i>associative arrays</i>, also called <i>dictionaries</i> in Python and <i>maps</i> in Scala and Java.</p> <p><u>In Spark, RDDs containing key-value tuples are called pair RDDs</u>. Although you don't have to use data in Spark in the form of key-value pairs (as you've seen in the previous chapters), pair RDDs are well suited (and indispensable) for many use cases. <u>Having keys along with the data enables you to aggregate, sort, and join the data</u>, as you'll soon see. But before doing any of that, the first step is to create a pair RDD, of course.</p>

Claim Chart – Claim 1

Claim Feature	Evidence from Charles Schwab
<p>partitioning the data of each one of the data groups into a plurality of data partitions that each have a plurality of key-value pairs and providing each data partition to a selected one of a plurality of mapping functions that are each user-configurable to independently output a plurality of lists of values for each of a set of keys found in such map function's corresponding data partition to form corresponding intermediate data for that data group and identifiable to that data group,</p>	<p>The RDDs support two types of data operations: transformations and actions.</p> <p>[17] RDD Programming Guide at 6:</p> <div data-bbox="653 638 1995 987"> <p>RDD Operations</p> <p>RDDs support two types of operations: <u>transformations</u>, which create a new dataset from an existing one, and <i>actions</i>, which return a value to the driver program after running a computation on the dataset. <u>For example, map is a transformation that passes each dataset element through a function and returns a new RDD representing the results.</u> On the other hand, reduce is an action that aggregates all the elements of the RDD using some function and returns the final result to the driver program (although there is also a parallel reduceByKey that returns a distributed dataset).</p> <p>All transformations in Spark are <i>lazy</i>, in that they do not compute their results right away. Instead, they just remember the transformations applied to some base dataset (e.g. a file). The transformations are only computed when an action requires a result to be returned to the driver program. This design enables Spark to run more efficiently. For example, we can realize that a dataset created through map will be used in a reduce and return only the result of the reduce to the driver, rather than the larger mapped dataset.</p> </div>

Claim Chart – Claim 1

Claim Feature

partitioning the data of each one of the data groups into a plurality of data partitions that each have a plurality of key-value pairs and **providing each data partition to a selected one of a plurality of mapping functions** that are each user-configurable to independently output a plurality of lists of values for each of a set of keys found in such map function's corresponding data partition to form corresponding intermediate data for that data group and identifiable to that data group,

Evidence from Charles Schwab

Examples of transformation functions:

[17] RDD Programming Guide at 11–12:

Transformations [↗](#)

The following table lists some of the common transformations supported by Spark. Refer to the RDD API doc ([Scala](#), [Java](#), [Python](#), [R](#)) and pair RDD functions doc ([Scala](#), [Java](#)) for details.

Transformation	Meaning
<code>map(func)</code>	Return a new distributed dataset formed by passing each element of the source through a function <i>func</i> .
<code>filter(func)</code>	Return a new dataset formed by selecting those elements of the source on which <i>func</i> returns true.
<code>flatMap(func)</code>	Similar to map, but each input item can be mapped to 0 or more output items (so <i>func</i> should return a Seq rather than a single item).
<code>mapPartitions(func)</code>	Similar to map, but runs separately on each partition (block) of the RDD, so <i>func</i> must be of type <code>Iterator<T> => Iterator<U></code> when running on an RDD of type T.
<code>mapPartitionsWithIndex(func)</code>	Similar to mapPartitions, but also provides <i>func</i> with an integer value representing the index of the partition, so <i>func</i> must be of type <code>(Int, Iterator<T>) => Iterator<U></code> when running on an RDD of type T.
<code>groupByKey(numPartitions)</code>	When called on a dataset of (K, V) pairs, returns a dataset of (K, Iterable<V>) pairs. Note: If you are grouping in order to perform an aggregation (such as a sum or average) over each key, using <code>reduceByKey</code> or <code>aggregateByKey</code> will yield much better performance. Note: By default, the level of parallelism in the output depends on the number of partitions of the parent RDD. You can pass an optional <code>numPartitions</code> argument to set a different number of tasks.
<code>reduceByKey(func, [numPartitions])</code>	When called on a dataset of (K, V) pairs, returns a dataset of (K, V) pairs where the values for each key are aggregated using the given reduce function <i>func</i> , which must be of type <code>(V,V) => V</code> . Like in <code>groupByKey</code> , the number of reduce tasks is configurable through an optional second argument.
<code>aggregateByKey(zeroValue)(seqOp, combOp, [numPartitions])</code>	When called on a dataset of (K, V) pairs, returns a dataset of (K, U) pairs where the values for each key are aggregated using the given combine functions and a neutral "zero" value. Allows an aggregated value type that is different than the input value type, while avoiding unnecessary allocations. Like in <code>groupByKey</code> , the number of reduce tasks is configurable through an optional second argument.

Claim Chart – Claim 1

Claim Feature	Evidence from Charles Schwab
<p>partitioning the data of each one of the data groups into a plurality of data partitions that each have a plurality of key-value pairs and providing each data partition to a selected one of a plurality of mapping functions that are each user-configurable to independently output a plurality of lists of values for each of a set of keys found in such map function's corresponding data partition to form corresponding intermediate data for that data group and identifiable to that data group,</p>	<p>The transformation operations include narrow and wide transformations. Wide transformations (also called “shuffle”) operate on data from multiple partitions. Wide transformations include aggregateByKey, reduceByKey, etc.</p> <p>[19] Spark RDD Transformations at 2–3:</p> <div data-bbox="741 654 1793 1208"> <p>RDD Transformation Types</p> <p>There are two types are transformations.</p> <p>Narrow Transformation</p> <p>Narrow transformations are the result of <u>map()</u> and <u>filter()</u> functions and these compute data that live on a single partition meaning there will not be any data movement between partitions to execute narrow transformations.</p> <p>Functions such as map(), mapPartition(), flatMap(), filter(), union() are some examples of narrow transformation</p> <p>Wider Transformation</p> <p>Wider transformations are the result of <u>groupByKey()</u> and <u>reduceByKey()</u> functions and <u>these compute data that live on many partitions</u> meaning there will be data movements between partitions to execute wider transformations. Since these shuffles the data, they also called shuffle transformations.</p> <p>Functions such as groupByKey(), aggregateByKey(), aggregate(), join(), repartition() are some examples of a wider transformations.</p> </div>

Claim Chart – Claim 1

Claim Feature

partitioning the data of each one of the data groups into a plurality of data partitions that each have a plurality of key-value pairs and **providing each data partition to a selected one of a plurality of mapping functions** that are each user-configurable to independently output a plurality of lists of values for each of a set of keys found in such map function's corresponding data partition to form corresponding intermediate data for that data group and identifiable to that data group,

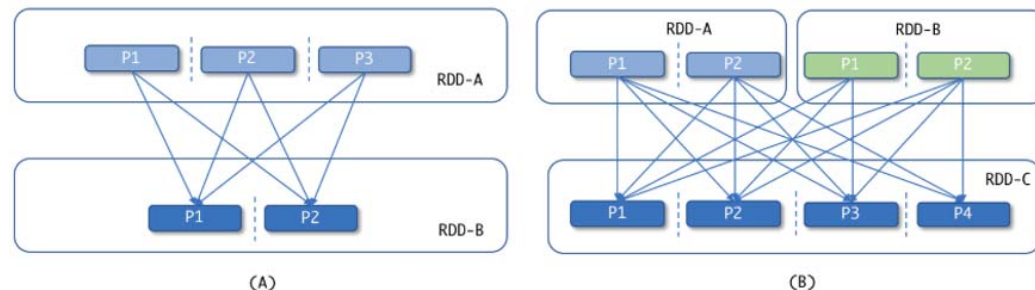
Evidence from Charles Schwab

[20]

Apache Spark Quick Start Guide:

Wide transformations

Wide transformations involve a shuffle of the data between the partitions. The `groupByKey()`, `reduceByKey()`, `join()`, `distinct()`, and `intersect()` are some examples of wide transformations. In the case of these transformations, the result will be computed using data from multiple partitions and thus requires a shuffle. Wide transformations are similar to the shuffle-and-sort phase of MapReduce. Let's understand the concept with the help of the following example:



Claim Chart – Claim 1

Claim Feature	Evidence from Charles Schwab
<p>partitioning the data of each one of the data groups into a plurality of data partitions that each have a plurality of key-value pairs and providing each data partition to a selected one of a plurality of mapping functions that are each user-configurable to independently output a plurality of lists of values for each of a set of keys found in such map function's corresponding data partition to form corresponding intermediate data for that data group and identifiable to that data group,</p>	<p>When executing a shuffle transformation, a map task is first run on all data partitions.</p> <p>[21] Spark SQL Shuffle Partitions:</p> <div data-bbox="1247 537 2018 1219"> <p>What is Spark Shuffle?</p> <p>Shuffling is a mechanism Spark uses to <u>redistribute the data</u> across different executors and even across machines. Spark shuffling triggers when we perform certain transformation operations like <code>groupByKey()</code>, <code>reduceByKey()</code>, <code>join()</code> on RDD and DataFrame.</p> <p>When <u>creating an RDD</u> or DataFrame, Spark doesn't necessarily store the data for all keys in a partition since at the time of creation there is no way we can set the key for data set.</p> <p>Hence, when we run the <code>reduceByKey()</code> operation to aggregate the data on keys, Spark does the following. needs to first run tasks to collect all the data from all partitions and</p> <p>For example, when we perform <code>reduceByKey()</code> operation, Spark does the following</p> <ul style="list-style-type: none"> • Spark first runs map tasks on all partitions which groups all values for a single key. • The results of the map tasks are kept in memory. • When results do not fit in memory, Spark stores the data into a disk. • Spark shuffles the mapped data across partitions, some times it also stores the shuffled data into a disk for reuse when it needs to recalculate. • Run the garbage collection • Finally runs reduce tasks on each partition based on key. </div>

Claim Chart – Claim 1

Claim Feature	Evidence from Charles Schwab
<p>partitioning the data of each one of the data groups into a plurality of data partitions that each have a plurality of key-value pairs and providing each data partition to a selected one of a plurality of mapping functions that are each user-configurable to independently output a plurality of lists of values for each of a set of keys found in such map function's corresponding data partition to form corresponding intermediate data for that data group and identifiable to that data group,</p>	<p>[22] Apache Spark Map vs. FlatMap at 3–4:</p> <div data-bbox="688 542 1955 922" style="border: 1px solid black; padding: 10px;"> <p>i. Spark Map Transformation</p> <p>A map is a transformation operation in Apache Spark. It applies to each element of RDD and it returns the result as new RDD. <u>In the Map, operation developer can define his own custom business logic.</u> The same logic will be applied to all the elements of RDD.</p> <p><u>Spark Map function takes one element as input process it according to custom code (specified by the developer) and returns one element at a time.</u> Map transforms an RDD of length N into another RDD of length N. The input and output RDDs will typically have the same number of records.</p> </div>

Claim Chart – Claim 1

Claim Feature	Evidence from Charles Schwab
<p>partitioning the data of each one of the data groups into a plurality of data partitions that each have a plurality of key-value pairs and providing each data partition to a selected one of a plurality of mapping functions that are each user-configurable to independently output a plurality of lists of values for each of a set of keys found in such map function's corresponding data partition to form corresponding intermediate data for that data group and identifiable to that data group,</p>	<p>When a mapping function is applied, intermediate files are created for that partition.</p> <p>[15] Spark in Action:</p> <div data-bbox="989 526 2018 1219"> <p>4.2.2. Understanding and avoiding unnecessary shuffling</p> <p>Physical movement of data between partitions is called <i>shuffling</i>. It occurs when data from multiple partitions needs to be combined in order to build partitions for a new RDD. When grouping elements by key, for example, <u>Spark needs to examine all of the RDD's partitions, find elements with the same key, and then physically group them, thus forming new partitions.</u></p> <p>To visualize what happens with partitions during a shuffle, we'll use the previous example with the <u>aggregateByKey transformation</u> (from section 4.1.2). The shuffle that occurs during that transformation is illustrated in <u>figure 4.2.</u></p> <p><u>The transform function puts all values of each key in a single partition (partitions P1 to P3) into a list. Spark then writes this data to intermediate files on each node. In the next phase, the merge function is called to merge lists from different partitions, but of the same key, into a single list for each key. The default partitioner (HashPartitioner) then kicks in and puts each key in its proper partition.</u></p> <p><u>Tasks that immediately precede and follow the shuffle are called <i>map</i> and <i>reduce</i> tasks, respectively. The results of <i>map</i> tasks are written to intermediate files (often to the OS's filesystem cache only) and read by reduce tasks. In addition to being written to disk, the data is sent over the network, so it's important to try to minimize the number of shuffles during Spark jobs.</u></p> </div>

Claim Chart – Claim 1

Claim Feature

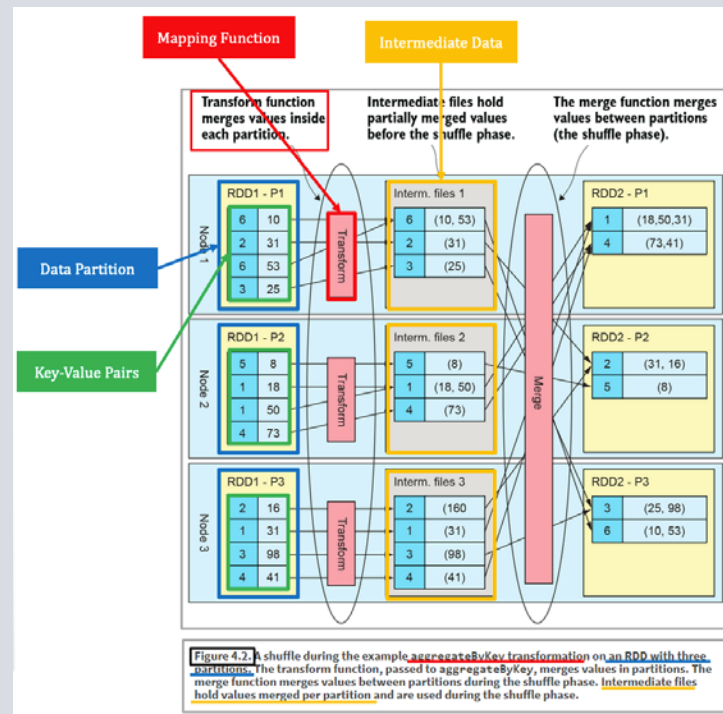
partitioning the data of each one of the data groups into a plurality of data partitions that each have a plurality of key-value pairs and providing each data partition to a selected one of a plurality of mapping functions that are each user-configurable to independently output a plurality of lists of values for each of a set of keys found in such map function's corresponding data partition to form corresponding intermediate data for that data group and identifiable to that data group,

Evidence from Charles Schwab

In this example, the aggregateByKey transformation operation (shuffle) is performed on three exemplary partitions:

[15]



Spark in Action:



Claim Chart – Claim 1

Claim Feature	Evidence from Charles Schwab
<p>wherein the data of a first data group has a different schema than the data of a second data group and the data of the first data group is mapped differently than the data of the second data group so that different lists of values are output for the corresponding different intermediate data, wherein the different schema and corresponding different intermediate data have a key in common; and</p>	<p>Users of the Charles Schwab System can create custom Mapper functions using Hive language and embed them into the MapReduce processing:</p> <p>[11] Hive Tutorial:</p> <div data-bbox="697 670 1860 1127" style="border: 1px solid black; padding: 10px;"> <p>Custom Map/Reduce Scripts</p> <p>Users can also plug in their own custom mappers and reducers in the data <u>stream by using features natively supported in the Hive language.</u> for example, in order to run a custom mapper script - map_script - and a custom reducer script - reduce_script - the user can issue the following command which uses the TRANSFORM clause to embed the mapper and the reducer scripts.</p> </div>

Claim Chart – Claim 1

Claim Feature	Evidence from Charles Schwab
<p>wherein the data of a first data group has a different schema than the data of a second data group and the data of the first data group is mapped differently than the data of the second data group so that different lists of values are output for the corresponding different intermediate data, wherein the different schema and corresponding different intermediate data have a key in common; and</p>	<p>Data partitions can be processed by using different user-configured Mappers:</p> <p>[4] HadoopMapReduce:</p> <div data-bbox="667 690 1917 1096" style="border: 1px solid black; padding: 10px;"> <p>As key-value pairs are read from the RecordReader they are passed to the configured  Mapper. The user supplied Mapper does whatever it wants with the input pair and calls  <u>OutputCollector.collect</u> with key-value pairs of its own choosing. The output it generates must use one key class and one value class. This is because the Map output will be written into a SequenceFile which has per-file type information and all the records must have the same type (use subclassing if you want to output different data structures). The Map input and output key-value pairs are not necessarily related typewise or in cardinality.</p> </div>

Claim Chart – Claim 1

Claim Feature	Evidence from Charles Schwab				
<p>wherein the data of a first data group has a different schema than the data of a second data group and the data of the first data group is mapped differently than the data of the second data group so that different lists of values are output for the corresponding different intermediate data, wherein the different schema and corresponding different intermediate data have a key in common; and</p>	<p>The Schema class in Hive is defined for each data set. Thus, the data of a first data group can have a different schema than the data of a second data group.</p> <p>[10] Class Schema:</p> <table border="1"> <thead> <tr> <th data-bbox="674 776 982 824">Modifier and Type</th><th data-bbox="1003 776 1350 824">Class and Description</th></tr> </thead> <tbody> <tr> <td data-bbox="674 889 924 938">static class</td><td data-bbox="1003 889 1944 1101"> <p><u>Schema.Fields</u></p> <p>The set of fields this struct contains, along with convenience methods for finding and manipulating them.</p> </td></tr> </tbody> </table>	Modifier and Type	Class and Description	static class	<p><u>Schema.Fields</u></p> <p>The set of fields this struct contains, along with convenience methods for finding and manipulating them.</p>
Modifier and Type	Class and Description				
static class	<p><u>Schema.Fields</u></p> <p>The set of fields this struct contains, along with convenience methods for finding and manipulating them.</p>				

Claim Chart – Claim 1

Claim Feature	Evidence from Charles Schwab
<p>wherein the data of a first data group has a different schema than the data of a second data group and the data of the first data group is mapped differently than the data of the second data group so that different lists of values are output for the corresponding different intermediate data, wherein the different schema and corresponding different intermediate data have a key in common; and</p>	<p>The Mapper function maps the key-value pairs to form intermediate key-value pairs. Similarly, different Mappers generate different intermediate results:</p> <p>[2] MapReduce Tutorial:</p> <div data-bbox="716 711 1881 1130" style="border: 1px solid black; padding: 10px; margin: 10px 0;"> <p>Mapper maps input key/value pairs to a set of intermediate key/value pairs.</p> <p><u>Maps are the individual tasks that transform input records into intermediate records.</u> The transformed intermediate records do not need to be of the same type as the input records. A given input pair may map to zero or many output pairs.</p> </div>

Claim Chart – Claim 1

Claim Feature	Evidence from Charles Schwab
<p>wherein the data of a first data group has a different schema than the data of a second data group and the data of the first data group is mapped differently than the data of the second data group so that different lists of values are output for the corresponding different intermediate data, wherein the different schema and corresponding different intermediate data have a key in common; and</p>	<p>A Hive JOIN query may be used to join two datasets that have a key in common:</p> <p>[12] LanguageManual Joins:</p> <div data-bbox="743 662 1852 1156" style="border: 1px solid black; padding: 10px;"> <p>Examples</p> <p>Some salient points to consider when writing join queries are as follows:</p> <ul style="list-style-type: none"> • Complex join expressions are allowed e.g. <pre>SELECT a.* FROM a JOIN b ON (a.id = b.id)</pre> <pre>SELECT a.* FROM a JOIN b ON (a.id = b.id AND a.department = b.department)</pre> <pre>SELECT a.* FROM a LEFT OUTER JOIN b ON (a.id <> b.id)</pre> </div>

Claim Chart – Claim 1

Claim Feature	Evidence from Charles Schwab
<p>wherein the data of a first data group has a different schema than the data of a second data group and the data of the first data group is mapped differently than the data of the second data group so that different lists of values are output for the corresponding different intermediate data, wherein the different schema and corresponding different intermediate data have a key in common; and</p>	<p>Spark supports multiple types of structure data, each with its own schema.</p> <p>[18] Learning Spark:</p> <div data-bbox="989 529 1885 781"> <p>Structured Data with Spark SQL</p> <p>Spark SQL is a component added in Spark 1.0 that is quickly becoming Spark's preferred way to work with structured and semistructured data. <u>By structured data, we mean data that has a <i>schema</i></u>—that is, a consistent set of fields across data records. <u>Spark SQL supports multiple structured data sources as input, and because it understands their schema</u>, it can efficiently read only the fields you require from these data sources. We will cover Spark SQL in more detail in Chapter 9, but for now, we show how to use it to load data from a few common sources.</p> </div> <div data-bbox="989 789 1885 927"> <p>Apache Hive</p> <p><u>One common structured data source on Hadoop is Apache Hive. Hive can store tables in a variety of formats, from plain text to column-oriented formats, inside HDFS or other storage systems.</u> Spark SQL can load any table supported by Hive.</p> </div> <div data-bbox="989 935 1885 1211"> <p>JSON</p> <p>If you have <u>JSON data with a consistent schema across records</u>, Spark SQL can infer their schema and load this data as rows as well, making it very simple to pull out the fields you need. To load JSON data, first create a HiveContext as when using Hive. (No installation of Hive is needed in this case, though—that is, you don't need a <i>hive-site.xml</i> file.) Then use the <code>HiveContext.jsonFile</code> method to get an RDD of Row objects for the whole file. Apart from using the whole Row object, you can also register this RDD as a table and select specific fields from it. For example, suppose that we had a JSON file containing tweets in the format shown in Example 5-33, one per line.</p> </div>

Claim Chart – Claim 1

Claim Feature	Evidence from Charles Schwab																					
wherein the data of a first data group has a different schema than the data of a second data group and the data of the first data group is mapped differently than the data of the second data group so that different lists of values are output for the corresponding different intermediate data, wherein the different schema and corresponding different intermediate data have a key in common; and	<div>[18] Learning Spark:</div> <div><div><div>File Formats</div><div>Spark makes it very simple to load and save data in a large number of file formats. Formats range from unstructured, like text, to semistructured, like JSON, to structured, like SequenceFiles (see Table 5-1). The input formats that Spark wraps all transparently handle compressed formats based on the file extension.</div></div><div><div>Table 5-1. Common supported file formats</div><table><tr><th>Format name</th><th>Structured</th><th>Comments</th></tr><tr><td>Text files</td><td>No</td><td>Plain old text files. Records are assumed to be one per line.</td></tr><tr><td>JSON</td><td>Semi</td><td>Common text-based format, semistructured; most libraries require one record per line.</td></tr><tr><td>CSV</td><td>Yes</td><td>Very common text-based format, often used with spreadsheet applications.</td></tr><tr><td>SequenceFiles</td><td>Yes</td><td>A common Hadoop file format used for key/value data.</td></tr><tr><td>Protocol buffers</td><td>Yes</td><td>A fast, space-efficient multilanguage format.</td></tr><tr><td>Object files</td><td>Yes</td><td>Useful for saving data from a Spark job to be consumed by shared code. Breaks if you change your classes, as it relies on Java Serialization.</td></tr></table></div></div>	Format name	Structured	Comments	Text files	No	Plain old text files. Records are assumed to be one per line.	JSON	Semi	Common text-based format, semistructured; most libraries require one record per line.	CSV	Yes	Very common text-based format, often used with spreadsheet applications.	SequenceFiles	Yes	A common Hadoop file format used for key/value data.	Protocol buffers	Yes	A fast, space-efficient multilanguage format.	Object files	Yes	Useful for saving data from a Spark job to be consumed by shared code. Breaks if you change your classes, as it relies on Java Serialization.
Format name	Structured	Comments																				
Text files	No	Plain old text files. Records are assumed to be one per line.																				
JSON	Semi	Common text-based format, semistructured; most libraries require one record per line.																				
CSV	Yes	Very common text-based format, often used with spreadsheet applications.																				
SequenceFiles	Yes	A common Hadoop file format used for key/value data.																				
Protocol buffers	Yes	A fast, space-efficient multilanguage format.																				
Object files	Yes	Useful for saving data from a Spark job to be consumed by shared code. Breaks if you change your classes, as it relies on Java Serialization.																				

Claim Chart – Claim 1

Claim Feature	Evidence from Charles Schwab
<p>wherein the data of a first data group has a different schema than the data of a second data group and the data of the first data group is mapped differently than the data of the second data group so that different lists of values are output for the corresponding different intermediate data, wherein the different schema and corresponding different intermediate data have a key in common; and</p>	<p>[16] Apache Spark FAQ at 1:</p> <p>How does Spark relate to Apache Hadoop?</p> <p>Spark is a fast and general processing engine compatible with Hadoop data. It can run in Hadoop clusters through YARN or Spark's standalone mode, and it can process data in HDFS, HBase, Cassandra, <u>Hive</u>, and any Hadoop InputFormat. It is designed to perform both batch processing (similar to MapReduce) and new workloads like streaming, interactive queries, and machine learning.</p> <p>[23] LanguageManual Transform at 1:</p> <p>Transform/Map-Reduce Syntax</p> <p><u>Users can also plug in their own custom mappers and reducers</u> in the data stream by using features natively supported in the <u>Hive</u> language. e.g. in order to run a custom mapper script - map_script - and a custom reducer script - reduce_script - the user can issue the following command which uses the TRANSFORM clause to embed the mapper and the reducer scripts.</p>

Claim Chart – Claim 1

Claim Feature

wherein the data of a first data group has a different schema than the data of a second data group and the data of the first data group is mapped differently than the data of the second data group so that different lists of values are output for the corresponding different intermediate data, **wherein the different schema and corresponding different intermediate data have a key in common;** and

Evidence from Charles Schwab

[15]

Spark in Action:

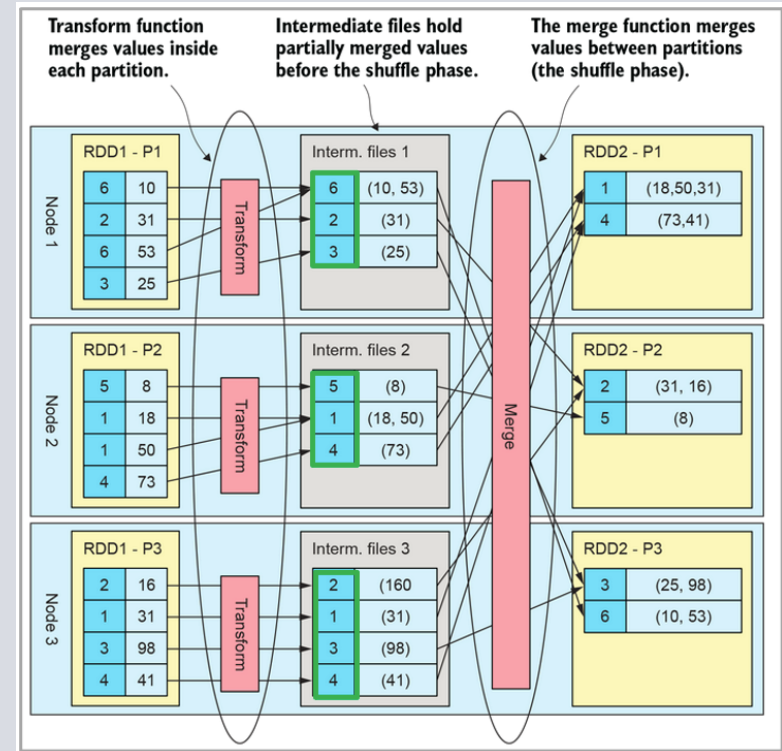


Figure 4.2. A shuffle during the example `aggregateByKey` transformation on an RDD with three partitions. The transform function, passed to `aggregateByKey`, merges values in partitions. The merge function merges values between partitions during the shuffle phase. Intermediate files hold values merged per partition and are used during the shuffle phase.

Claim Chart – Claim 1

Claim Feature	Evidence from Charles Schwab
<p>reducing the intermediate data for the data groups to at least one output data group, including processing the intermediate data for each data group in a manner that is defined to correspond to that data group, so as to result in a merging of the corresponding different intermediate data based on the key in common,</p>	<p>The method performed by the Charles Schwab system includes reducing the intermediate data for the data groups to at least one output data group, including processing the intermediate data for each data group in a manner that is defined to correspond to that data group, so as to result in a merging of the corresponding different intermediate data based on the key in common.</p> <p>The Reducer function reduces the intermediate values into a smaller set of values. This function operates in three phases: (1) Shuffle; (2) Sort; and (3) Reduce. The Shuffle and Sort phases operate on the output of various Mappers and group data before the final Reduce phase based on a common key.</p> <div data-bbox="577 836 1932 1218"> <p>[2] MapReduce Tutorial:</p> <p>Reducer</p> <p><u>Reducer reduces a set of intermediate values which share a key to a smaller set of values.</u></p> <p>The number of reduces for the job is set by the user via <code>Job.setNumReduceTasks(int)</code>.</p> <p>Overall, Reducer implementations are passed the Job for the job via the <code>Job.setReducerClass(Class)</code> method and can override it to initialize themselves. The framework then calls <code>reduce(WritableComparable, Iterable<Writable>, Context)</code> method for each <key, (list of values)> pair in the grouped inputs. Applications can then override the <code>cleanup(Context)</code> method to perform any required cleanup.</p> </div>

Claim Chart – Claim 1

Claim Feature	Evidence from Charles Schwab
<p>reducing the intermediate data for the data groups to at least one output data group, including processing the intermediate data for each data group in a manner that is defined to correspond to that data group, so as to result in a merging of the corresponding different intermediate data based on the key in common,</p>	<p>[3] Interface Reducer:</p> <p>Reduces a set of intermediate values which share a key to a smaller set of values.</p> <p>The number of Reducers for the job is set by the user via <code>JobConf.setNumReduceTasks(int)</code>. Reducer implementations can access the <code>JobConf</code> for the job via the <code>JobConfigurable.configure(JobConf)</code> method and initialize themselves. Similarly they can use the <code>Closeable.close()</code> method for de-initialization.</p> <p>Reducer has 3 primary phases:</p> <ol style="list-style-type: none"> 1. Shuffle <p>Reducer is input the grouped output of a Mapper. In the phase the framework, for each Reducer, fetches the relevant partition of the output of all the Mappers, via HTTP.</p> 2. Sort <p>The framework groups Reducer inputs by keys (since different Mappers may have output the same key) in this stage.</p> <p>The shuffle and sort phases occur simultaneously i.e. while outputs are being fetched they are merged.</p> <p>Secondary Sort</p> <p>If equivalence rules for keys while grouping the intermediates are different from those for grouping keys before reduction, then one may specify a <code>Comparator</code> via <code>JobConf.setOutputValueGroupingComparator(Class)</code>. Since <code>JobConf.setOutputKeyComparatorClass(Class)</code> can be used to control how intermediate keys are grouped, these can be used in conjunction to simulate <i>secondary sort on values</i>.</p> <p>For example, say that you want to find duplicate web pages and tag them all with the url of the "best" known example. You would set up the job like:</p> <ul style="list-style-type: none"> • Map Input Key: url • Map Input Value: document • Map Output Key: document checksum, url pagerank • Map Output Value: url • Partitioner: by checksum • OutputKeyComparator: by checksum and then decreasing pagerank • OutputValueGroupingComparator: by checksum <ol style="list-style-type: none"> 3. Reduce <p>In this phase the <code>reduce(Object, Iterator, OutputCollector, Reporter)</code> method is called for each <key, (list of values)> pair in the grouped inputs.</p>

Claim Chart – Claim 1

Claim Feature	Evidence from Charles Schwab
<p>reducing the intermediate data for the data groups to at least one output data group, including processing the intermediate data for each data group in a manner that is defined to correspond to that data group, so as to result in a merging of the corresponding different intermediate data based on the key in common,</p>	<p>The Reduce function operates on different sets of input data (intermediate data) that have a key in common. From this processing, the Reduce function populates an output file.</p> <p>[3] Interface Reducer:</p> <div data-bbox="1142 638 2028 1222" style="border: 1px solid black; padding: 10px;"> <p>reduce</p> <pre>void reduce(K2 key, Iterator<V2> values, OutputCollector<K3,V3> output, Reporter reporter) throws IOException</pre> <p><u>Reduces values for a given key.</u></p> <p>The framework calls this method for each <key, (list of values)> pair in the grouped inputs. Output values must be of the same type as input values. Input keys must not be altered. The framework will reuse the key and value objects that are passed into the reduce, therefore the application should clone the objects they want to keep a copy of. In many cases, all values are combined into zero or one value.</p> <p><u>Output pairs are collected with calls to <code>OutputCollector.collect(Object, Object)</code>.</u></p> <p>Applications can use the Reporter provided to report progress or just indicate that they are alive. In scenarios where the application takes a significant amount of time to process individual key/value pairs, this is crucial since the framework might assume that the task has timed-out and kill that task. The other way of avoiding this is to set <code>mapreduce.task.timeout</code> to a high-enough value (or even zero for no time-outs).</p> <p>Parameters:</p> <p>key - the key.</p> <p>values - the list of values to reduce.</p> <p>output - to collect keys and combined values.</p> <p>reporter - facility to report progress.</p> <p>Throws:</p> <p>IOException</p> </div>

Claim Chart – Claim 1

Claim Feature	Evidence from Charles Schwab
<p>reducing the intermediate data for the data groups to at least one output data group, including processing the intermediate data for each data group in a manner that is defined to correspond to that data group, so as to result in a merging of the corresponding different intermediate data based on the key in common,</p>	<p>In Spark, the intermediate data is reduced to an output data group.</p> <p>[21] Spark SQL Shuffle Partitions:</p> <div data-bbox="1178 537 1969 1230" style="border: 1px solid black; padding: 10px;"> <p>What is Spark Shuffle?</p> <p>Shuffling is a mechanism Spark uses to <u>redistribute the data</u> across different executors and even across machines. Spark shuffling triggers when we perform certain transformation operations like <code>groupByKey()</code>, <code>reduceByKey()</code>, <code>join()</code> on RDD and DataFrame.</p> <p>When <u>creating an RDD</u> or DataFrame, Spark doesn't necessarily store the data for all keys in a partition since at the time of creation there is no way we can set the key for data set.</p> <p>Hence, when we run the <code>reduceByKey()</code> operation to aggregate the data on keys, Spark does the following. needs to first run tasks to collect all the data from all partitions and</p> <p>For example, when we perform <code>reduceByKey()</code> operation, Spark does the following</p> <ul style="list-style-type: none"> • Spark first runs <code>map</code> tasks on all partitions which groups all values for a single key. • The results of the map tasks are kept in memory. • When results do not fit in memory, Spark stores the data into a disk. • Spark shuffles the mapped data across partitions, some times it also stores the shuffled data into a disk for reuse when it needs to recalculate. • Run the garbage collection • Finally runs reduce tasks on each partition <u>based on key.</u> </div>

Claim Chart – Claim 1

Claim Feature	Evidence from Charles Schwab
<p>reducing the intermediate data for the data groups to at least one output data group, including processing the intermediate data for each data group in a manner that is defined to correspond to that data group, so as to result in a merging of the corresponding different intermediate data based on the key in common,</p>	<p>[15] Spark in Action:</p> <div data-bbox="989 500 2018 1198"> <p>4.2.2. Understanding and avoiding unnecessary shuffling</p> <p>Physical movement of data between partitions is called <i>shuffling</i>. It occurs when data from multiple partitions needs to be combined in order to build partitions for a new RDD. When grouping elements by key, for example, Spark needs to examine all of the RDD's partitions, find elements with the same key, and then physically group them, thus forming new partitions.</p> <p>To visualize what happens with partitions during a shuffle, we'll use the previous example with the <u>aggregateByKey transformation</u> (from section 4.1.2). The shuffle that occurs during that transformation is illustrated in <u>figure 4.2</u>.</p> <p>The transform function puts all values of each key in a single partition (partitions P1 to P3) into a list. Spark then writes this data to intermediate files on each node. <u>In the next phase, the merge function is called to merge lists from different partitions, but of the same key, into a single list for each key.</u> The default partitioner (HashPartitioner) then kicks in and puts each key in its proper partition.</p> <p><u>Tasks that immediately precede and follow the shuffle are called map and reduce tasks, respectively.</u> <u>The results of map tasks are written to intermediate files (often to the OS's filesystem cache only) and read by reduce tasks.</u> In addition to being written to disk, the data is sent over the network, so it's important to try to minimize the number of shuffles during Spark jobs.</p> </div>

Claim Chart – Claim 1

Claim Feature

reducing the intermediate data for the data groups to at least one output data group, including processing the intermediate data for each data group in a manner that is defined to correspond to that data group, so as to result in a merging of the corresponding different intermediate data based on the key in common,

Evidence from Charles Schwab

[15]

Spark in Action:

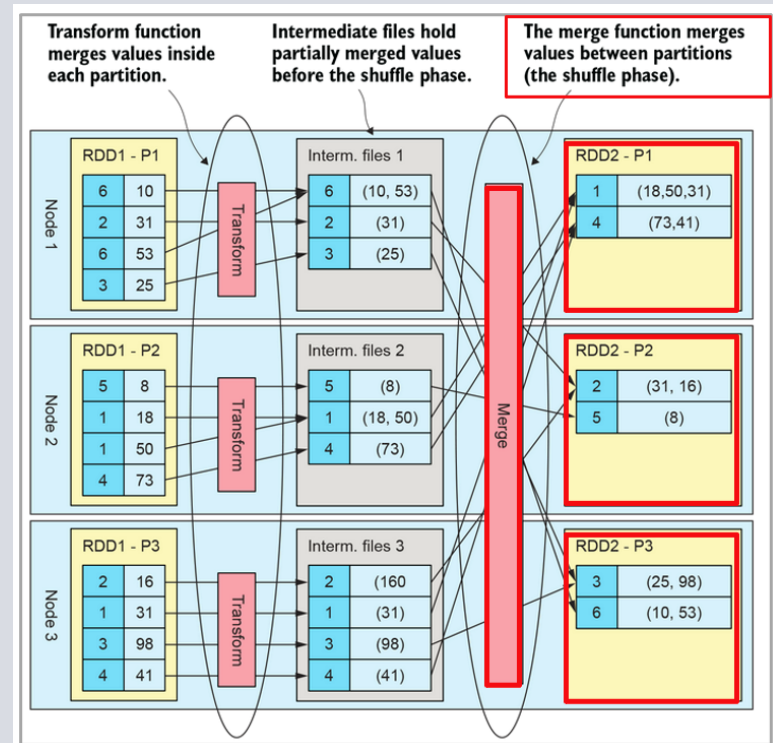


Figure 4.2: A shuffle during the example `aggregateByKey` transformation on an RDD with three partitions. The transform function, passed to `aggregateByKey`, merges values in partitions. The merge function merges values between partitions during the shuffle phase. Intermediate files hold values merged per partition and are used during the shuffle phase.

Claim Chart – Claim 1

Claim Feature	Evidence from Charles Schwab
<p>wherein the mapping and reducing operations are performed by a distributed system.</p>	<p>The mapping and reducing operations are performed by the Charles Schwab system, which is a distributed system.</p> <p>[1] An introduction to Apache Hadoop for big data:</p> <div data-bbox="1050 641 2026 1193" style="border: 1px solid black; padding: 10px;"> <p>The Apache Hadoop framework is composed of the following modules</p> <ol style="list-style-type: none"> 1. Hadoop Common: contains libraries and utilities needed by other Hadoop modules 2. <u>Hadoop Distributed File System (HDFS): a distributed file-system that stores data on the commodity machines, providing very high aggregate bandwidth across the cluster</u> 3. <u>Hadoop YARN: a resource-management platform responsible for managing compute resources in clusters and using them for scheduling of users' applications</u> 4. <u>Hadoop MapReduce: a programming model for large scale data processing</u> <p><u>All the modules in Hadoop are designed with a fundamental assumption that hardware failures (of individual machines, or racks of machines) are common and thus should be automatically handled in software by the framework.</u> Apache Hadoop's MapReduce</p> </div>

END

